Harvey Product Vision: Data Architecture, and the Agentic Future

Building the Collaborative Bridge: From Law Firm Innovation to Shared Client Intelligence.

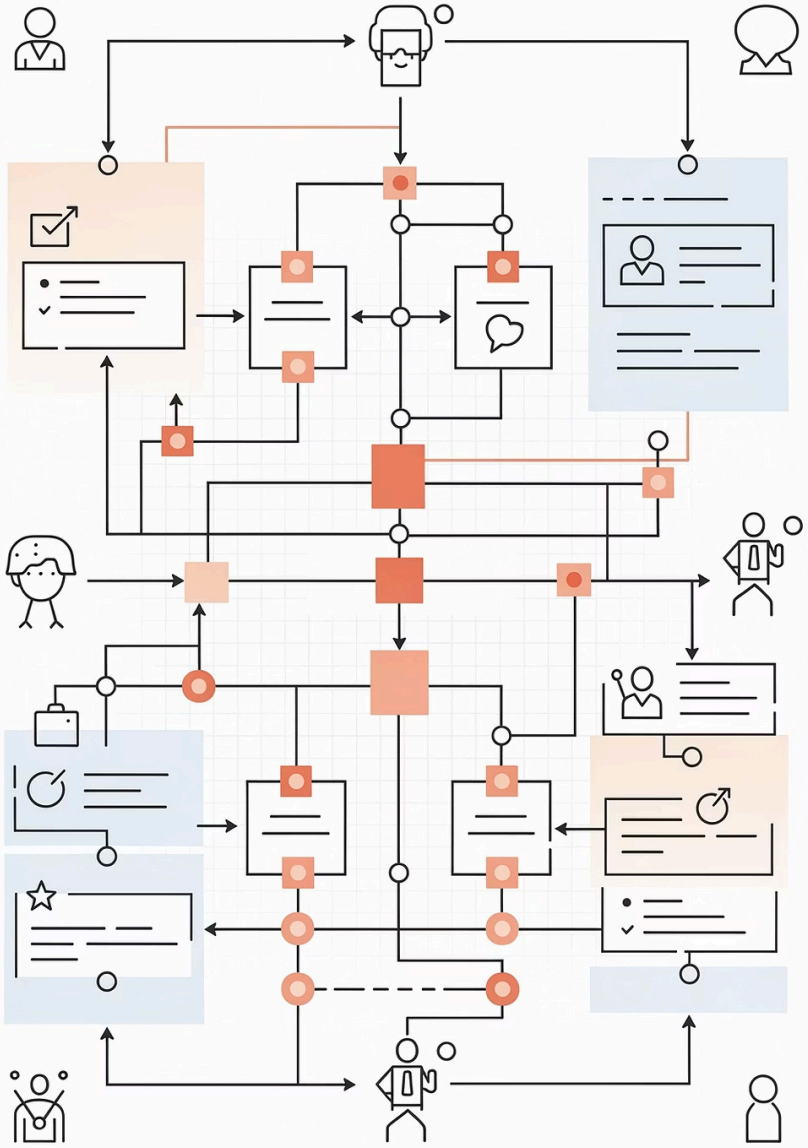

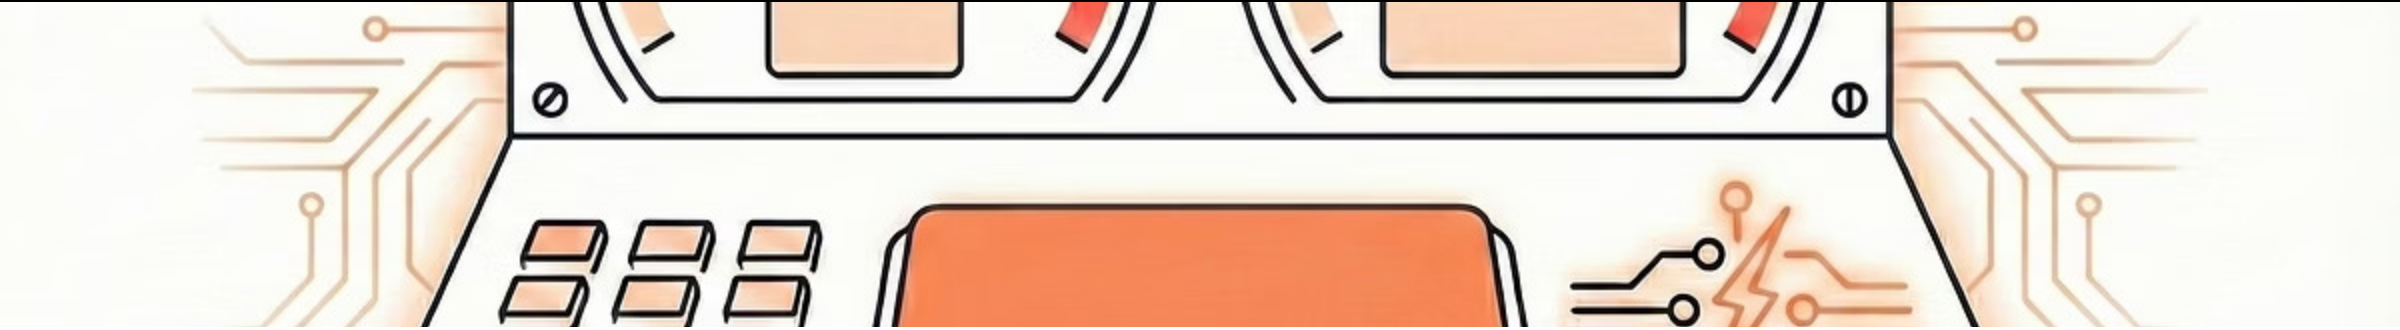

The Core Philosophy: Re-Earning the Edge

The 'Wrapper' Risk: Any feature that is merely a UI wrapper will be absorbed by the next base model update.

"Wrapper" Divergence:

- Model-Agnosticism: Harvey leverages multiple AI providers (OpenAI, Anthropic, and Gemini) to avoid vendor lock-in and optimize for different use cases
- The Infrastructure Sacrifice: Harvey committed to a 6-month focus on building security infrastructure over UI features to establish defensibility in regulated industries

The Vertical Layer (Harvey)

Harvey builds three critical vertical layers on top of generic models:

1. Complex Permissioning

Advanced access controls for legal data

2. Ethical Walls

Prevent data leakage between customers

3. Specialized Legal Data Handling

Domain-specific processing & compliance

To avoid the 'wrapper' risk (generic models absorbing UI-only features), Harvey builds defensible, specialized security infrastructure for regulated industries, even if it means short-term speed sacrifices.

> "You need to re-earn your position as a non-wrapper company every maybe like three months." — Winston Weinberg

Executive Summary: From LLM Wrapper to Enterprise Infrastructure

Three key strategic shifts:

1. Information to Intelligence

Harvey is moving from simple text processing (redlines) to capturing the underlying legal logic. The core value asset is the 'Decision Trace'—the reasoning behind the edit, not just the edit itself.

2. The 'Scaffolding' Strategy

Prioritizing 'Banking-Grade' security over rapid UI iteration. This involves building deep permissioning layers, auditability, and ethical walls to support highly regulated client data.

3. The Agentic Shift

Preparing for a future of 'Long Horizon' tasks. Moving from chatbots that answer in seconds to autonomous agents that work unmonitored for days to execute complex workflows.

ROI: Capturing Market Share

AI allows law firms to take on large-scale 'mundane' projects they previously rejected due to resource constraints. By automating high-volume, repetitive work, firms can transform these projects into strategic revenue drivers while maintaining profitability and freeing senior talent for higher-value work.

The Data Engine: Capturing 'Decision Traces'

The 'Decision Trace' captures the metadata of legal logic—the reasoning, fallback clauses, and negotiation strategy embedded in comment bubbles—not just the final text change. This allows Harvey to learn why decisions were made, enabling the system to replicate a legal team's judgment rather than simply mimicking text patterns.

Traditional Data

Focuses on the final text of a contract clause and changes (redlines). This offers limited learning value.

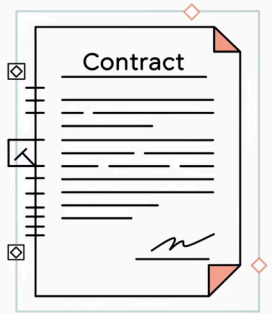

Decision Trace

Captures reasoning, fallback clauses, and negotiation leverage within comment bubbles. This enables the system to replicate a legal team's judgment.

Traditional LLMs like ChatGPT primarily use chat history. Legal work relies on complex documents. Harvey transforms document metadata into memory artifacts, effectively solving context window limitations by understanding the underlying logic, not just the text.

Product Innovation: Document-Based Memory

The four-step process:

Ingestion

Weekly rolling analysis of uploaded legal documents.

Extraction

Isolating 'Decision Traces' and metadata from the files.

Vector Storage

Artifacts stored in Vector Database for semantic retrieval.

Retrieval

System retrieves logic from past work, not just text matches.

Quality Assurance

The Weekly Rolling Review

Human-in-the-loop step where lawyers audit and prune the AI's memory map, ensuring accuracy and relevance of captured decision traces.

Generic LLMs (ChatGPT) rely on chat history. Legal work relies on documents. Harvey converts document metadata into memory artifacts to solve context window limitations.

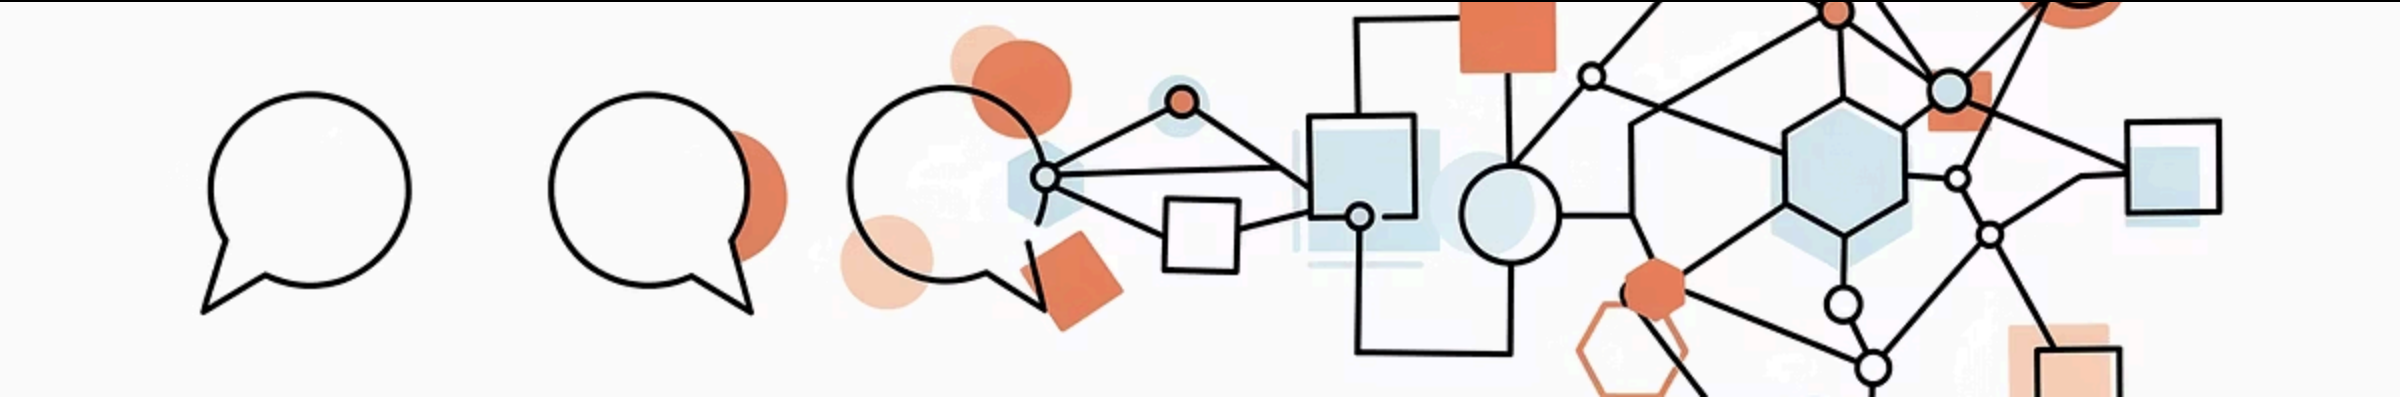

The Future of Automation: Agentic Systems

Current State (Chatbot)

Short Horizon Task: "Summarize this PDF"

Prompt → Immediate Response

Future State (Agentic System - 2026)

Long Horizon Task (3 Days): "Analyze these 50 contracts against our playbook, flag risks, and draft a summary."

Prompt → Autonomous Analysis → Cross-referencing → Risk Flagging → Drafting → Final Output

The Prerequisite

Agents operating unmonitored for days require pre-built trust in the permissioning architecture (The Scaffolding).

Implementation: The 'IA before AI'

"The in-house teams are like about a year behind I'd say... most of them now are interested in like piloting the tools and checking them out, but change management across a lot of in-house teams [hasn't] happened."

— **Winston Weinberg**, Co-Founder & CEO, Harvey

AI cannot have 'Memory' if documents are unstructured or tagged as 'Other'. Buying the license is easy; cleaning the data is the hard work.

> "One CIO cannot do the work of 20 people... you have to invest in the teams that manage this."

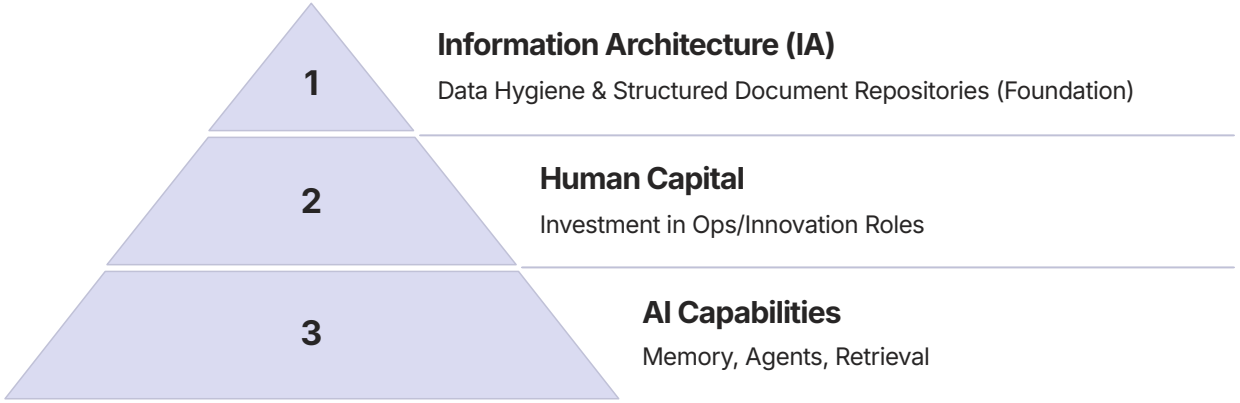

The Human Catalyst

- **Partner Buy-In:** The #1 predictor of AI implementation success. Without leadership commitment and active participation, even the best technology will fail to gain traction.
- **Hiring for Slope:** Finding talent that scales with AI evolution. Look for individuals who can adapt, learn, and grow alongside rapidly changing AI capabilities rather than those with fixed skill sets.

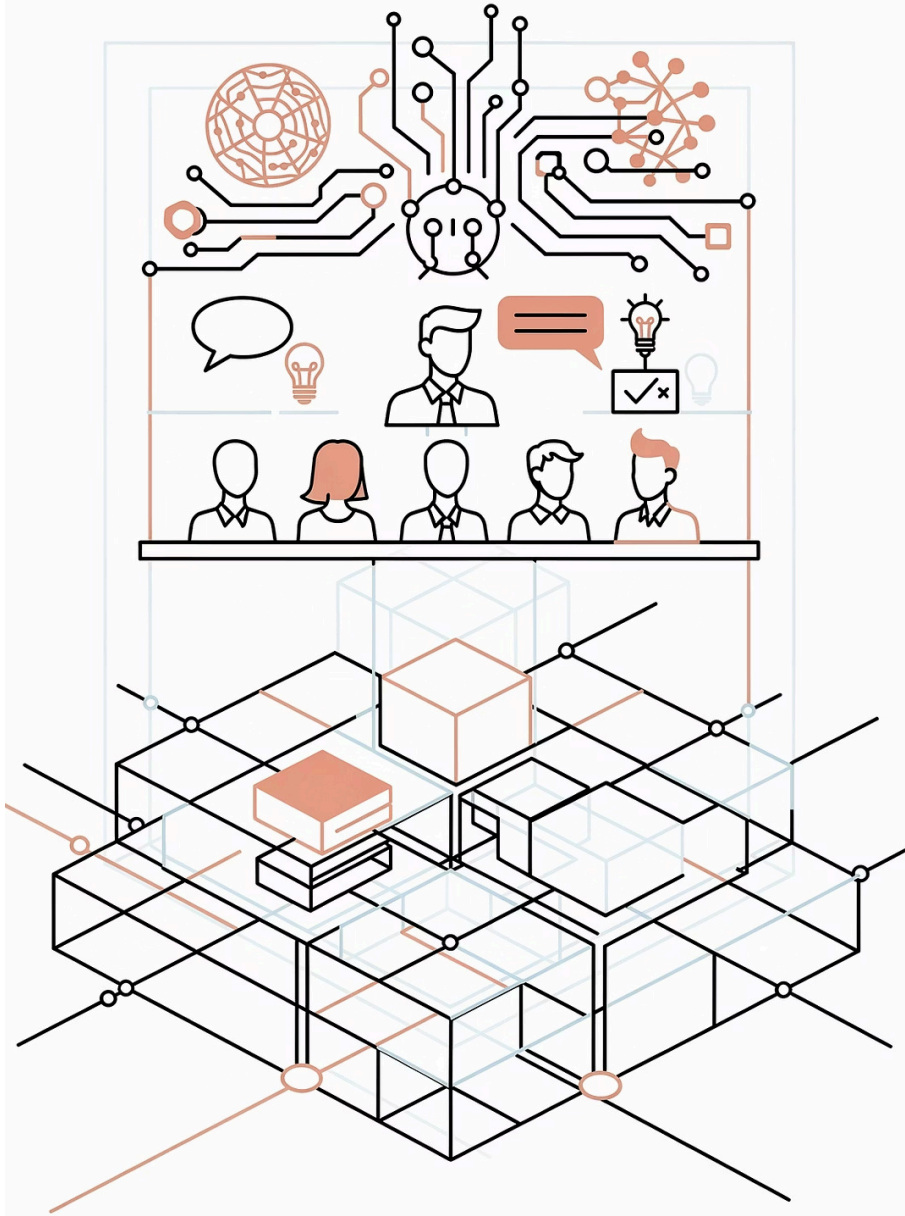

Appendix: Technical Lexicon & Concepts

Decision Traces

Metadata of legal work (comments, logic, reasoning) extracted to train the system on why a decision was made, not just what was changed.

Vector Database

A storage mechanism that allows the AI to retrieve documents based on semantic meaning and concept rather than simple keyword matching.

Agentic Systems

AI capable of performing "Long Horizon" tasks (hours/days) autonomously, executing multi-step workflows rather than single-prompt responses.

Scaffolding

The underlying security, permissioning, and ethical wall architecture required to safely deploy autonomous agents in regulated environments.

Shared Spaces

Collaborative environments where law firms and clients work together within the same platform, enabling real-time knowledge sharing and joint decision-making while maintaining appropriate security boundaries.

Model Agnosticism

Strategic approach of leveraging multiple LLM providers (OpenAI, Anthropic, Gemini) to optimize for different use cases, avoid vendor lock-in, and ensure the best model is used for each specialized task.